Integrative Systems

Firing of Putative Dopamine Neurons in Ventral Tegmental Area Is Modulated by Probability of Success during Performance of a Stop-Change Task

Stephen S. Tennyson,[1] Adam T. Brockett,[1,2] Nicholas W. Hricz,[1] Daniel W. Bryden,[1,2] and Matthew R. Roesch[1,2]

[1]Department of Psychology, University of Maryland, College Park, MD 20742 and [2]Program in Neuroscience and Cognitive Science, University of Maryland, College Park, MD 20742

Abstract

Response inhibition, the ability to refrain from unwanted actions, is an essential component of complex behavior and is often impaired across numerous neuropsychiatric disorders such as addiction, attention-deficit hyperactivity disorder (ADHD), schizophrenia, and obsessive-compulsive disorder. Accordingly, much research has been devoted to characterizing brain regions responsible for the regulation of response inhibition. The stop-signal task, a task in which animals are required to inhibit a prepotent response in the presence of a STOP cue, is one of the most well-studied tasks of response inhibition. While pharmacological evidence suggests that dopamine (DA) contributes to the regulation of response inhibition, what is exactly encoded by DA neurons during performance of response inhibition tasks is unknown. To address this issue, we recorded from single units in the ventral tegmental area (VTA), while rats performed a stop-change task. We found that putative DA neurons fired less and higher to cues and reward on STOP trials relative to GO trials, respectively, and that firing was reduced during errors. These results suggest that DA neurons in VTA encode the uncertainty associated with the probability of obtaining reward on difficult trials instead of the saliency associated with STOP cues or the need to resolve conflict between competing responses during response inhibition.

Key words: conflict; dopamine; inhibition; neuron; rat; stop signal

---

Significance Statement

The ability to refrain from unwanted actions, also known as response inhibition, is an essential component of complex behavior, and is impaired across numerous neuropsychiatric disorders, including addiction, attention-deficit hyperactivity disorder (ADHD), and schizophrenia. Dopamine (DA) is important for reward learning, but its role in response inhibition is less clear. For the first time, we characterized the activity of DA neurons in rats performing a response inhibition task and found that DA neurons primarily signaled information regarding the uncertainty of obtaining reward during cues and reward delivery when behavioral trials were difficult and there was a low probability of success.

Introduction

The ability to resolve conflict between competing responses and inhibit unwanted actions, also known as cognitive control, is an essential component of complex

behavior. Cognitive control is impaired in numerous neuropsychiatric disorders, including schizophrenia (Bellgrove et al., 2006; Dajani and Uddin, 2015), attention-deficit hyperactivity disorder (ADHD; Oosterlaan and

Sergeant, 1998; Oosterlaan et al., 1998; Durston et al., 2009), and substance abuse disorders (Fillmore and Rush, 2002; Monterosso et al., 2005). The wide array of symptoms associated with these phenotypically distinct disorders highlights the importance of cognitive control in daily life, but also suggests that research regarding the neural mechanisms supporting conflict detection may provide useful insights into the pathologic etiology of these disorders.

Across species, and in clinical populations, a common paradigm used to test cognitive control and response inhibition is the stop-signal task (Verbruggen and Logan, 2008; Eagle and Baunez, 2010; Boecker et al., 2013). During performance of the stop-signal task, participants respond quickly to a "GO" cue (e.g., light or tone) by performing an instrumental response (e.g., button press, lever press, etc.). On GO trials (~80% of all trials), participants develop an automatic tendency to respond quickly to the presentation of the GO cue. On "STOP" trials (~20% of all trials), participants must inhibit this prepotent GO response when the STOP cue is presented (e.g., second light or tone). Difficulty arises from the automaticity induced by the high proportion of GO trials leading to decreased accuracy on STOP trials. In "stop-change" variants of this task, participants are not only required to inhibit their behavior on STOP trials, but also to redirect their behavior in the opposite direction (Bryden et al., 2011, 2012, 2018; Boecker et al., 2013; Bryden and Roesch, 2015). While much work has gone into the development and characterization of these tasks using pharmacological techniques, much less is known about the neural underpinnings that support this behavior.

Dopamine (DA) plays an essential role in reinforcement learning and decision-making (Roesch et al., 2007; Schultz 2013; Wood et al., 2017); however, its role in response inhibition has been incompletely studied. The use of drugs that target the DA system have yielded conflicting results ranging from improved performance on STOP trials, to altered performance on GO trials (Tannock et al., 1989; Aron et al., 2003; Bedard et al., 2003; Boonstra et al., 2005; Lijffijt et al., 2006; Eagle et al., 2008; Eagle and Baunez, 2010). Based on these results, it is difficult to parse the exact role DA plays in modulating performance on stop-signal tasks, because it remains unclear what is signaled by DA neurons during STOP tasks. It is known that separate populations of DA neurons can signal either changes in value associated with reward prediction errors or changes in the saliency of the cue, independent of its value (Matsumoto and Hikosaka 2009; Bromberg-Martin et al., 2010). That is, some DA neurons have been observed to fire more strongly for cues that predict a higher probability reward (vs cues that predict low probability reward) and also fire more strongly to delivery of unlikely reward (i.e., prediction error encoding), whereas other DA neurons increase firing to salient or alerting events independent from their value, thought to be critical for orienting and executive control (Bromberg-Martin et al., 2010).

During performance of stop-signal tasks, the activity of DA neurons may reflect reward prediction error encoding, such that DA neurons may fire less to STOP cues because they predict lower probability of reward, and fire more to successful reward delivery on STOP trials because reward delivery was less common. Alternatively, DA neurons might fire strongly to STOP cues due to their salient unexpected appearance. To test these possibilities, we recorded from putative DA neurons as rats performed our stop-change task. We found that overall DA firing was higher on GO trials during the response period, but higher on STOP trials at the time of reward. Moreover, we show that trials during which the rat was delayed in inhibiting and redirecting its behavior (i.e., response conflict), DA firing during the presentation of the STOP cue was reduced and firing during reward was more pronounced compared to trials during which response conflict was resolved more quickly. Finally, we show a correlation between activity and probability of success on difficult STOP trials such that firing was reduced after STOP cues that were preceded by multiple GO trials. Overall, these data suggest that DA firing in the ventral tegmental area (VTA) reflects the low probability of receiving a reward on STOP trials rather than a need to inhibit behavior on STOP trials or the salience associated with the low occurrence of STOP cues.

Materials and Methods

Animals

Four male and three female Long-Evans rats ($n = 7$; weight, 175–200 g) were obtained from Charles River Laboratories. Rats were housed on a 12/12 h light/dark schedule and all behavioral testing and recordings occurred between 9 A.M. and 2 P.M. All studies were approved by the Institutional Animal Care and Use Committee and conformed to the National Research Council Guide of the Care and Use of Laboratory Animals (2011).

Surgical procedures and histology

Surgical procedures followed guidelines for aseptic technique. Electrodes were manufactured and implanted as in prior recording experiments (Bryden et al., 2011, 2012; Bryden and Roesch, 2015). Rats were chronically implanted with a drivable bundle of 10 25 μm in diameter FeNiCr wires (Stablohm 675, California Fine Wire) in the VTA, counterbalanced across left and right hemispheres. Four animals were implanted at 5.2 mm posterior to bregma, 0.7 mm laterally, and 7.0 mm ventral to the brain surface as in prior experiments (Roesch et al., 2007), the remaining three animals were implanted with a 5° angle pointed at the midline, with coordinates at 5.2 mm posterior to bregma, 1.4 mm laterally, and 7.5 mm ventral to

the brain surface. Immediately before implantation, wires were freshly cut with surgical scissors to extend ~1 mm beyond the cannula and electroplated with platinum (H_2PtCl_6; Aldrich) to an impedance of ~300 kOhm. Cephalexin (15 mg kg^{-1}, postoperative) was administered twice daily for two weeks postoperatively. After recording, rats were perfused and their brains removed and processed for histology (Roesch et al., 2006).

Stop-change task

Recording was conducted in aluminum chambers ~18" on each side with downward sloping walls narrowing to an area of 12" × 12" at the bottom. On one wall, a central port was located above two adjacent fluid wells. Two directional lights were located above the two fluid wells. House lights were located above the panel. Task control was implemented via computer. Port entry, licking, and well entry times were monitored by disruption of photobeams.

The basic trial design is illustrated in Figure 1A,B. Each trial began by illumination of house lights that instructed the rat to nose poke into the central port. Nose poking initiated a 1000 ms pre-cue delay period. At the end of this delay, a directional light to the animal's left or right was flashed for 100 ms. If the rat exited the port at any time before offset of the directional cue light, the trial was aborted and house lights were extinguished. On 80% of trials, presentation of the left or right light signaled the direction in which the animal could respond to obtain sucrose reward in the corresponding fluid well below. On 20% of trials, the light opposite to the location of the originally cued direction turned on either at the same time as port exit or after a stop-signal delay (0–100 ms) and remained illuminated until the behavioral response was made. These trials will be referred to as STOP trials, which were randomly interleaved with GO trials. Rats were required to stop the movement signaled by the first light and respond in the direction of the second light. On correct responding, rats were required to remain in the fluid well for a variable period between 800 and 1000 ms (pre-fluid delay) before reward delivery (10% sucrose solution). Error trials (incorrect direction) were immediately followed by the extinction of house lights and ITI onset of 4 s. Trials were presented in a pseudorandom sequence such that left and right trials were presented in equal numbers (±1 over 250 trials).

Single-unit recordings

Procedures were the same as described previously (Bryden and Roesch, 2015). Wires were screened for activity daily; if no activity was detected, the rat was removed and the electrode assembly was advanced 40 or 80 μm. Otherwise, a session was conducted, and the electrode was advanced at the end of the session. Neural activity was recorded using four identical Plexon Multichannel Acquisition Processor Systems. Signals from electrode wires were amplified 20× by an op-amp headstage located on the electrode array. Immediately outside the training chamber, signals were passed through a differential pre-amplifier (Plexon Inc, PBX2/16sp-r-G50/16fp-G50) where single unit signals were amplified 50×

and filtered at 150–9000 Hz. The single unit signals were then sent to the Multichannel Acquisition Processor box, where they were further filtered at 250–8000 Hz, digitized at 40 kHz and amplified at 1–32x. Waveforms (>2.5:1 signal-to-noise) were extracted from active channels and recorded to disk by an associated workstation with event timestamps from the behavior computer.

DA cell identification

Neurons were screened for wide wave form and small amplitude characteristics using MATLAB as in prior experiments (Roesch et al., 2007; Takahashi et al., 2009, 2011; Takahashi and Schoenbaum, 2016; Jo et al., 2013; Sadacca et al., 2016; Park and Moghaddam, 2017; Wood et al., 2017). Wave form half duration and amplitude ratio (negative minus positive peak/sum) were calculated and clustered using k-means (Roesch et al., 2007; Takahashi et al., 2009, 2011; Takahashi and Schoenbaum, 2016; Sadacca et al., 2016). The center and variance of each cluster was computed without data from the neuron of interest, and then that neuron was assigned to a cluster if it fell within 3 SDs from the center of that cluster. If a neuron met the criteria for more than one cluster, it was not classified. This process was repeated for all neurons. Cells that increased firing to reward delivery (1 s; Wilcoxon, $p < 0.05$) and fell in the cluster with the longest half duration and smallest amplitude ratio were considered putative DA neurons (Roesch et al., 2007; Takahashi et al., 2009, 2011; Jo et al., 2013; Sadacca et al., 2016; Takahashi and Schoenbaum, 2016; Park and Moghaddam, 2017; Wood et al., 2017). We recorded from 809 VTA neurons from seven rats (1: 46 cells; 2: 141 cells; 3: 110 cells; 4: 137 cells; 5: 158 cells; 6: 120 cells; 7: 97 cells) during performance of a stop-change task (Fig. 1A,B), 85 neurons were classified as being putative DA. Of those 85 putative DA neurons, 77 were also responsive during the response epoch (i.e., nose poke exit to well entry; Wilcoxon, $p < 0.05$). A total of 475 neurons were classified as non-DA. Neurons that did not fall into one cluster or the other were excluded from analysis.

Data analysis

Units were sorted via Offline Sorter software from Plexon Inc, using a template matching algorithm, and analyzed in Neuroexplorer and MATLAB. Activity was examined during the period between nose poke exit and well entry (response epoch), the 800 ms period following well entry (post-response epoch), and the 500 ms period following reward delivery (reward epoch). Activity in population histograms was normalized by dividing by the maximal firing rate of each neuron. Activity was averaged across direction (e.g., responding left or right) given that DA neurons are not directionally selective (Roesch et al., 2007; Wood et al., 2017). All statistical procedures were executed using raw firing rates. Unless otherwise specified, behavioral data were analyzed using a two-way ANOVA, where each datum is a session average.
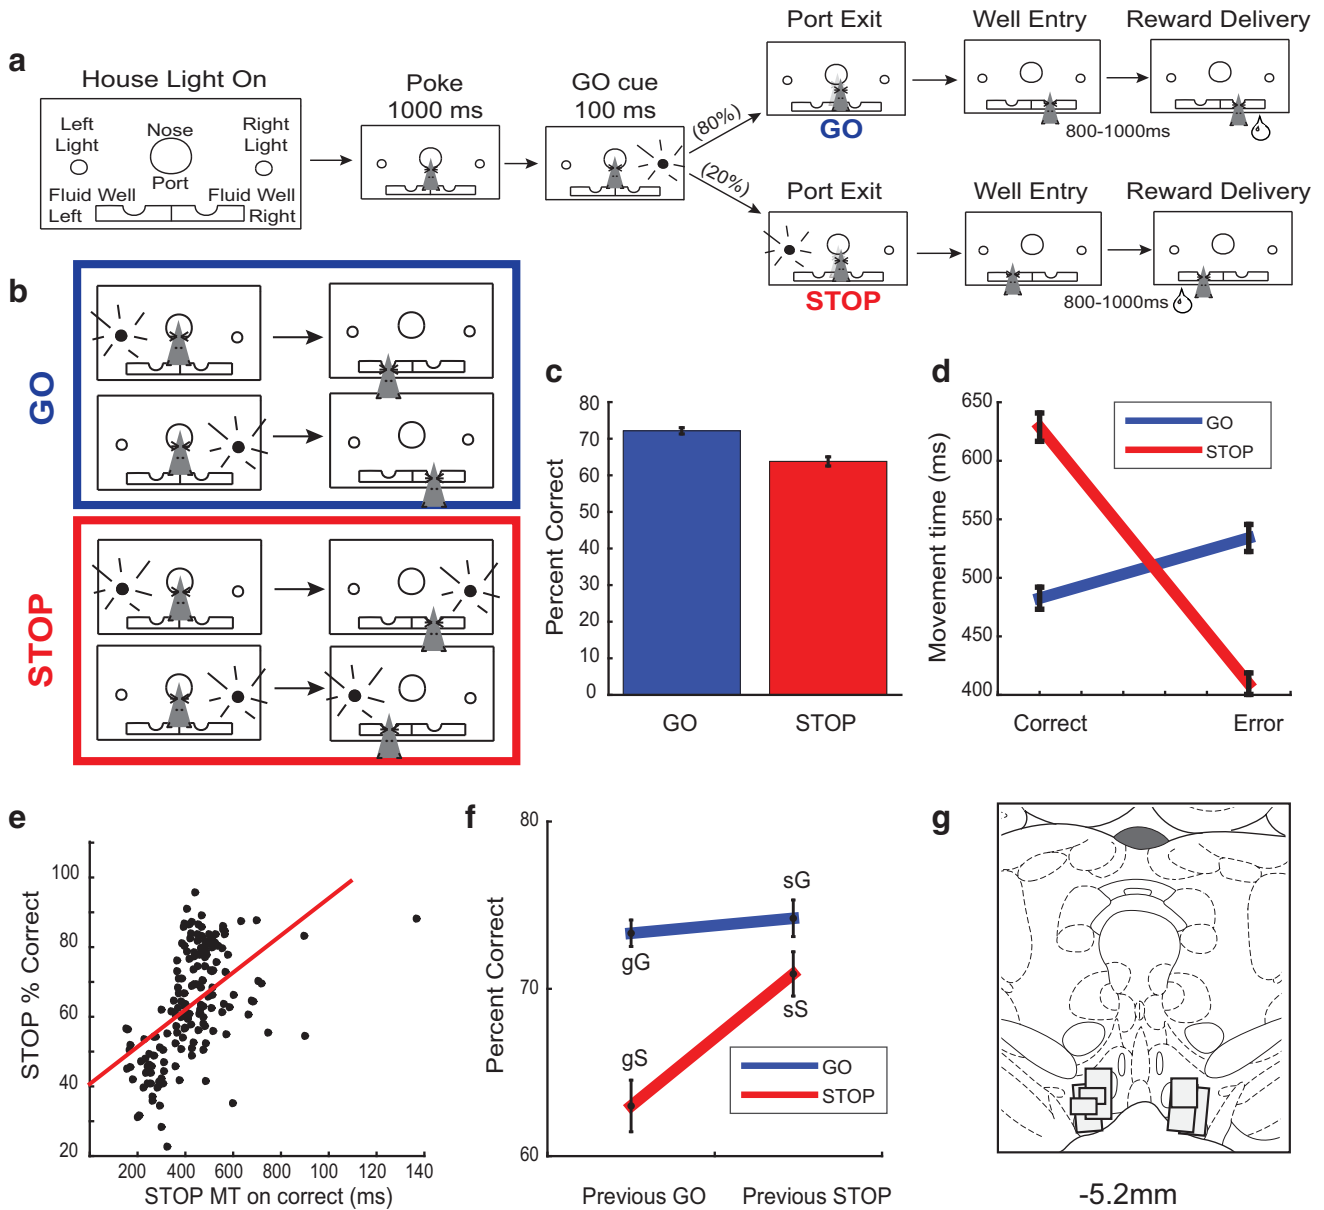
Figure 1. Task design and behavioral analysis. **A**, House lights signaled the rat to nose poke and wait inside a central odor port for 1000 ms before one of two directional cue lights were illuminated for 100 ms, directing the rat to the left or right adjacent fluid well to receive fluid reward. On 20% of trials, on port exit, the opposite-sided cue light would illuminate, requiring the rat to stop the initial action and respond in the opposite direction to receive reward. After entering the correct fluid well, rats were required to hold in the well for a variable period between 800 and 1000 ms before reward delivery. **B**, There are three trial types used in our analysis, STOP trials, GO trials, and STOP errors by two directions for each. **C**, Bar graph shows percentage correct scores as a function of all trials. **D**, Movement time behavior across correct and incorrect STOP (red) and GO (blue) trials. **E**, Correlation between STOP movement time (in seconds) on correct STOP trials versus average STOP percentage correct across each session (t test, $p < 0.05$). **F**, The impact of previous trial on current trial performance and conflict adaptation. Percentage correct shown across GO and STOP trials preceded by either a GO (gG, gS, left column) or STOP trial (sG, sS, right column). **G**, Location of recording sites (Paxinos and Watson, 2007). Boxes mark the extent of recording locations based on histology.

Results

Behavior

Rats were trained to respond to left and right cue lights that directed behavior to fluid wells for reward. Our analyses examined behavior averaged over sessions rather than averaged across sessions within each rat and then across rats. Such an analysis better represents the average behavior that occurs during collection of single neuron activity that will be presented below. The behavior in this task has been replicated in several studies in rats performing the same task (Bryden et al., 2011, 2012, 2018; Bryden and Roesch, 2015). Animals exhibited significantly reduced accuracy on STOP trials compared to

GO trials (*t* test: $t_{(175)} = 6.44$, $p < 0.001$; Fig. 1*C*). Rats were also significantly slower on STOP trials (Fig. 1*D*); a two-way ANOVA revealed a significant main effect of correctness ($F_{(1,643)} = 62.07$, $p < 0.001$) and a significant interaction of correctness by trial type ($F_{(1,643)} = 161.67$, $p < 0.001$). We observed no significant main effect for trial type ($F_{(1,643)} = 1.04$, $p = 0.31$). Movement times on STOP trial errors were faster than correct STOP trial types, indicating that animals failed to inhibit the initial GO response (Fig. 1*D*). Finally, rats' performance exhibited a speed-accuracy trade-off, in that when they were slower they tended to perform better on STOP trials (Pearson's correlation; $r = 0.52$; $p < 0.001$; Fig. 1*E*).

Performance on the current trial depended on the difficulty of the previous trial type. To determine the effects of previous trial type on the current trial's performance, percentage correct was analyzed across all possible combinations of current and previous trials [i.e., when STOP preceded STOP (sS); GO preceded STOP (gS); GO preceded GO (gG); STOP preceded GO (sG)]. STOP trials after GO trials were more difficult as evidenced by worse performance (Fig. 1*F*). A two-way ANOVA revealed a significant main effect of the previous trial type ($F_{(1,651)} = 13.01$, $p < 0.001$), a significant main effect of the current trial type ($F_{(1,651)} = 31.40$, $p < 0.001$), and a significant interaction between the previous and current trial type ($F_{(1,651)} = 8.26$, $p = 0.004$), demonstrating that percentage correct was significantly lower on STOP trials preceded by GO trials and that rats demonstrated conflict adaptation such that they were more accurate on STOP trials immediately following STOP trials.

Putative DA neurons show lower firing to STOP cues during the response period but higher for STOP trials at the time of reward

We recorded from 809 VTA neurons from seven rats (1: 46 cells; 2: 141 cells; 3: 110 cells; 4: 137 cells; 5: 158 cells; 6: 120 cells; 7: 97 cells) during performance of a stop-change task (Fig. 1*A,B*), 85 of which were classified as being putative DA (see methods). The recording locations are illustrated in Figure 1*G*. We hypothesized that DA firing would reflect reward prediction error encoding, such that DA neurons would fire less to STOP cues, but more to STOP rewards. Consistent with this hypothesis, we found that many putative DA neurons fired more for GO cues over STOP cues during the response epoch, and more for STOP rewards over GO rewards during the reward epoch. This is illustrated in Figure 2*A–D*, which displays average firing of a single putative DA neuron aligned to port exit and reward delivery. The example neuron showed weaker firing on STOP trials after port exit during illumination of the STOP cue (Fig. 2*A*), and stronger firing on STOP trials at the time of reward delivery (Fig. 2*B–D*).

To determine whether DA neurons fired differently on STOP versus GO trials, we averaged firing rate across the 85 putative DA neurons and aligned activity to the initial GO cue (Fig. 2*E*), port exit (Fig. 2*F*), well entry (Fig. 2*G*), and reward delivery (Fig. 2*H*). During initial GO cue pre-

sentation, firing rate increased non-distinctly across all three trial types (Fig. 2*E*). After port exit, which is the time when the STOP cue was illuminated on STOP trials, firing appeared higher on GO (blue) trials compared to STOP (red) trials (Fig. 2*F*). We generated a trial-type index (STOP – GO/STOP + GO) for firing rates taken from the time of port exit to well entry (response epoch) on correct trials to determine whether the firing rate significantly differed between STOP and GO trials (Wilcoxon, $p < 0.05$). We found that the distribution of trial-type indices was significantly shifted below zero (Wilcoxon, $\mu = -0.02$; $p = 0.01$; df = 84; Fig. 2*I*) and the counts of neurons that fired significantly more on GO trials outnumber those with the opposite effect (black bars; 3 STOP vs 11 GO; $\chi^2 = 4.46$; $p = 0.03$; Fig. 2*I*), suggesting that firing was stronger during the response epoch on GO versus STOP trials at both the population and single unit level.

Overall, firing appeared lower on STOP errors compared to correct STOP trials during the period following well entry (Fig. 2*G*). To quantify this effect, we computed an error index (STOP error – STOP correct/STOP error + STOP correct) on firing rates taken from the post-response epoch (800 ms after entering the fluid well) to determine whether the firing rate significantly differed between STOP correct and error trials across the entire population (Wilcoxon, $p < 0.05$). We found that the distribution was significantly shifted below zero (Wilcoxon, $\mu = -0.08$; $p = 0.009$; df = 84; Fig. 2*J*) and the counts of neurons that fired significantly more on correct STOP trials outnumbers those that fired more on STOP errors (black bars; 8 STOP error vs 19 STOP correct; $\chi^2 = 4.40$; $p = 0.04$; Fig. 2*J*), suggesting that putative DA firing was stronger on correct STOP trials versus incorrect STOP trials after well entry.

Finally, we asked whether firing was higher during STOP trials when reward was delivered relative to GO trials by examining firing aligned to reward delivery (Fig. 2*H*). For this analysis, we excluded error trials from the alignment because reward is not delivered during incorrect trials. We found that firing was only slightly stronger on STOP compared to GO trials during the time of reward delivery. To quantify this effect, we computed the trial-type index (STOP – GO/STOP + GO) during the 500 ms period after reward was delivered (reward epoch). The distribution of indices was significantly shifted about zero, indicating that the number of neurons that fired more on STOP than GO trials were in the majority (Wilcoxon, $\mu = 0.02$; $p = 0.02$; df = 84; Fig. 2*K*). Despite the significant positive shift in the population distribution, the counts of neurons that fired significantly more on STOP trials within a session was not greater than those that fired more on GO trials at the time of reward (black bars; 3 STOP vs 1 GO; $\chi^2 = 0.90$; $p = 0.34$; Fig. 2*K*). Overall these findings suggest that putative DA firing was modestly stronger on STOP versus GO trials at the time of reward delivery, suggesting that reward delivery after successful completion of a STOP trial elicited higher firing compared to rewards delivered after correct GO trials.
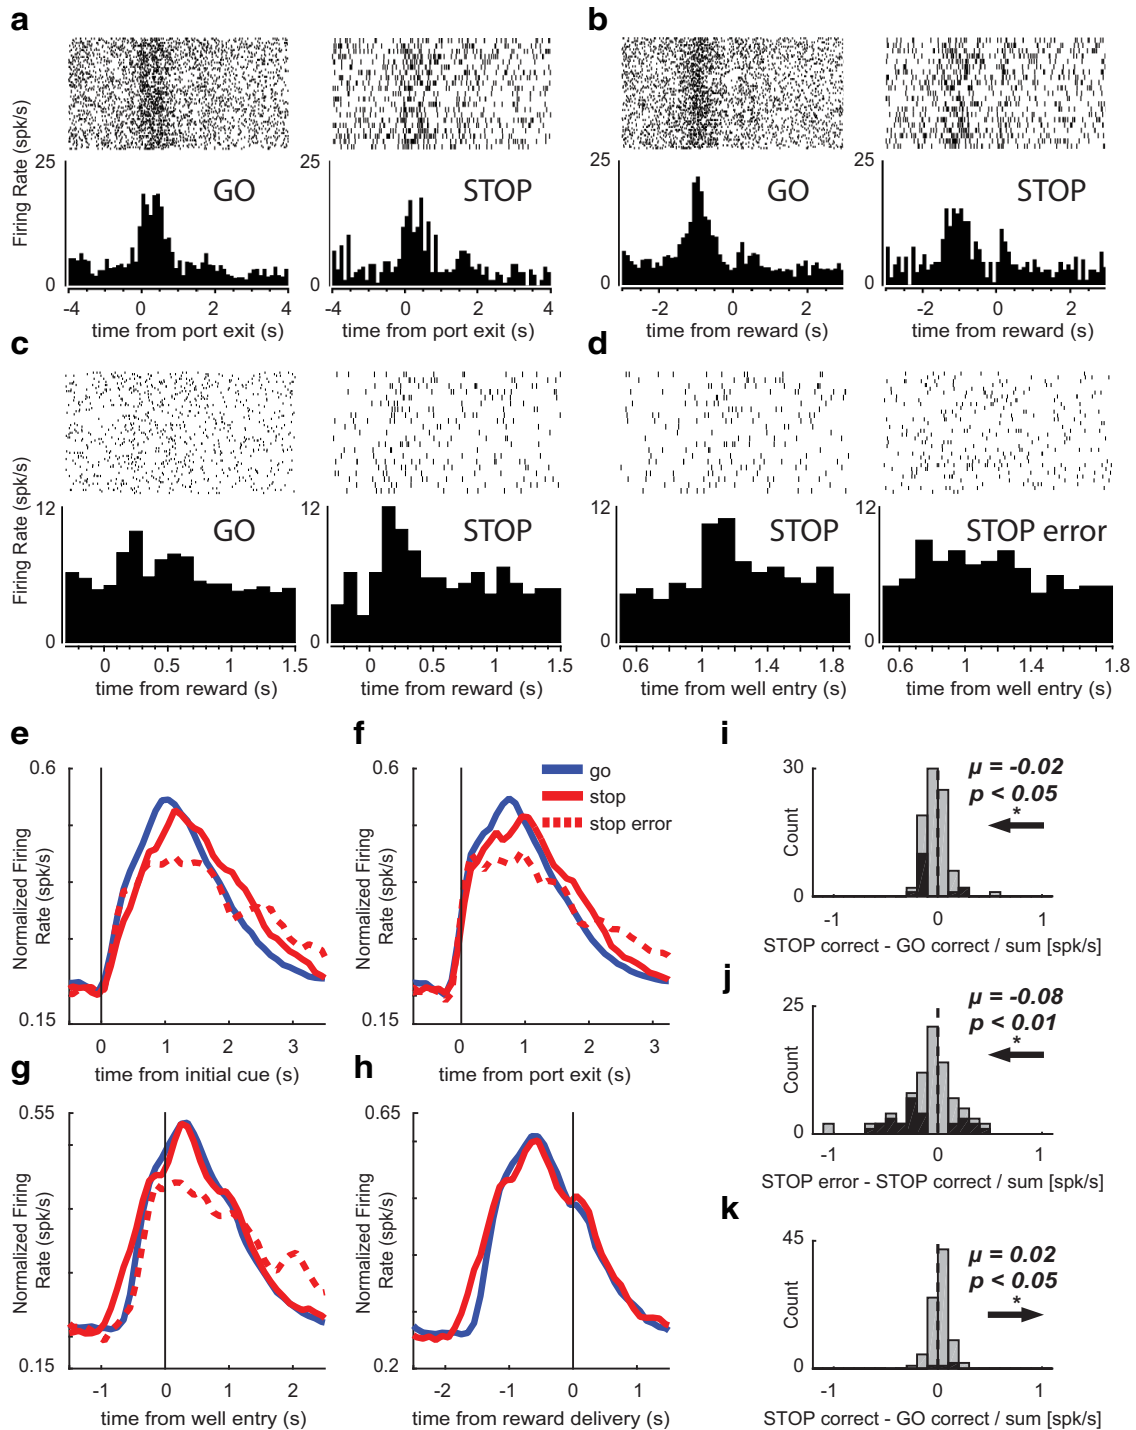
Figure 2. Putative DA firing was higher for GO trials during the response period but higher for STOPs at the time of reward. *A*, Single-cell example of a putative DA neuron aligned to port exit that fired more for GO trials (left) than STOP trials (right). *B*, The same single-cell example from *A* aligned to reward delivery, showing higher firing to reward on STOP trials (right) compared to GO trials (left). *C*, The same single-cell example from *A*, *B* aligned to reward delivery and zoomed in on the reward response, showing higher firing on STOP trials (right) compared to GO trials (left). *D*, The same single-cell example from *A*, *B* zoomed in on the response following well entry, showing higher firing on correct STOP trials (left) compared to STOP errors (right). *E–H*, Population average histograms of putative DA neurons plotted over trial time for GO (blue), STOP (red), and STOP error (dashed) trials while aligning to multiple events (n = 85). Activity is aligned to initial GO cue (*E*), port exit (*F*), well entry (*G*), and reward delivery (*H*). *E*, At the time of the initial cue, firing rate increases rapidly for all trial types. *F*, During the response epoch (the period from port exit to well entry, when the STOP cue was illuminated on STOP trials), firing was significantly higher on GO trials compared to STOP trials. *G*, During the post-response epoch (the 800 ms period following well entry), firing was significantly lower on STOP errors compared to STOP corrects. *H*, During the reward epoch (500 ms period after reward delivery), firing was significantly higher on STOP trials than GO trials.

Error trials are excluded since reward is not delivered on errors. *I–K*, Population average distributions for significant effects described above. Arrows depict direction of distribution shift for significant effects. Black bars represent the number of neurons that showed a significant difference between GO and STOP corrects (*I, K*) and correct and incorrect STOP trials (*J*; Wilcoxon, $p < 0.05$). Distributions are determined to be significantly different from zero via Wilcoxon. *I, K*, Firing rate during the response and reward epoch were compared during correct STOPs and GOs by computing a trial type index (STOP correct – GO correct/STOP correct + GO correct). *J*, Firing rate during the post-response epoch was compared in correct and incorrect STOP trials (STOP error – STOP correct/STOP error + STOP correct).

Putative DA neurons fire less to STOP cues, but more for STOP rewards, when the rat responded more slowly

The degree of conflict associated with making the appropriate response varies from trial to trial during a session. One measure of how difficult it is to resolve conflict on any given trial is to determine how long rats take to successfully perform a STOP trial. That is, the more difficult the trial, the longer it takes a rat to inhibit and redirect behavior. To determine whether DA activity was modulated by the speed with which animals responded, average population histograms were split into fast and slow trials based on movement times within each session. To determine whether firing rate was significantly different between fast and slow trial types, we calculated speed indices on firing rates to compare fast and slow GO trials (GO fast – GO slow), fast and slow STOP trials (STOP fast – STOP slow), and fast and slow STOP errors (fast STOP error – slow STOP error). As before, we examined differences in firing across three behavioral epochs (response epoch, post-response epoch, and reward epoch).

Firing appeared higher on fast STOP trials compared to slow STOP trials during the response epoch (Fig. 3*A,B*). We found no significant differences in firing rate between fast and slow GO trials (Wilcoxon, $\mu = -0.02$; $p = 0.46$; df = 84; Fig. 3*E*), or fast and slow STOP errors at the time of port exit (Wilcoxon, $\mu = -0.29$; $p = 0.12$; df = 84; Fig. 3*G*). However, putative DA neurons fired significantly less on slow STOP trials compared to fast STOP trials after port exit (response epoch; Wilcoxon, $\mu = 0.48$; $p = 0.02$; df = 84; Fig. 3*F*).

During the post-response epoch, firing appeared lower on fast STOP errors compared to slow STOP errors (Fig. 3*C,D*). After well entry, no significant differences in firing were apparent between fast and slow GO trials (Wilcoxon, $\mu = -0.19$; $p = 0.11$; df = 84; Fig. 3*H*) or fast and slow STOP trials (Wilcoxon, $\mu = -0.16$; $p = 0.27$; df = 84; Fig. 3*I*) were found. However, firing rates were significantly lower on fast STOP errors versus slow STOP errors during the post-response epoch (Wilcoxon, $\mu = -1.27$; $p < 0.001$; df = 84; Fig. 3*J*), suggesting that putative DA neurons fired more on slower compared to faster STOP errors.

Lastly, we examined putative DA population aligned to reward delivery, where firing appeared to be higher on slow STOP rewards compared to fast STOP rewards (Fig. 3*K,L*). We found no significant differences in firing between fast and slow GO trials (Wilcoxon, $\mu = -0.16$; $p = 0.30$; df = 84; Fig. 3*M*); however, DA neurons did fire significantly more during slower STOP trials compared to faster STOP trials (Wilcoxon, $\mu = -0.38$; $p = 0.04$; df = 84; Fig. 3*N*).

Putative DA neuron firing was modulated by heightened response conflict induced by previous trial type

We investigated whether changes in difficulty induced by the previous trial modulated putative DA firing during performance on the current trial. Recall that rats perform better on STOP trials that followed a STOP trial (i.e., conflict adaptation; Fig. 1*F*). To determine whether the DA signal was impacted by the modulation of behavior due to the previous trial type, we examined average activity plotted on correct GO trials, STOP trials preceded by a single GO trial (gS), and STOP trials preceded by a STOP trial (sS). The average firing rate over time is illustrated to in Figure 4*A,C*. As described above, the average firing rate was higher on GO compared to STOP trials, but we found little difference between sS (orange) and gS (red) trials.

To quantify these effects, we computed indices on firing rates to compare gS to GO trials (gS – GO/gS + GO), sS to GO trials (sS – GO/sS + GO), and gS to sS trials (gS – sS/gS + sS) across the three behavioral epochs (response epoch, post-response epoch, and reward epoch). During the response epoch, firing on GO trials was significantly higher than both gS and sS trials (GO vs gS: Wilcoxon, $\mu = -0.04$, $p = 0.01$, df = 84; GO vs sS: Wilcoxon, $\mu = -0.03$, $p = 0.01$, df = 84). These findings demonstrate that putative DA firing is higher on low conflict GO trials compared to either gS or sS trials during the response epoch as described above; however, firing rates between sS and gS trials were not significantly different from each other in any of the analysis epochs (response epoch: Wilcoxon, $\mu = -0.004$, $p = 0.71$, df = 84; post-response epoch: Wilcoxon, $\mu = -0.01$; $p = 0.43$; reward epoch: Wilcoxon, $\mu = 0.01$; $p = 0.93$) indicating that DA firing was not modulated on STOP trials by the nature of the previous trial (i.e., STOP or GO).

The lack of difference between sS and gS trials might reflect lack of encoding by DA neurons for this aspect of the task or that differences in difficulty between the two trial types was not strong enough to elicit differences in neural responding. To address this issue, we extended our analyses to study the effect that a train of multiple uninterrupted GO trials have on DA firing and percentage correct on STOP trials. Theoretically, the more GO trials that precede a STOP trial, the more difficult it would be to inhibit the response, thus lowering the probability of success. Indeed, we found a negative correlation between the number of previous GO trials and accuracy on the current STOP trial, such that STOP trial performance
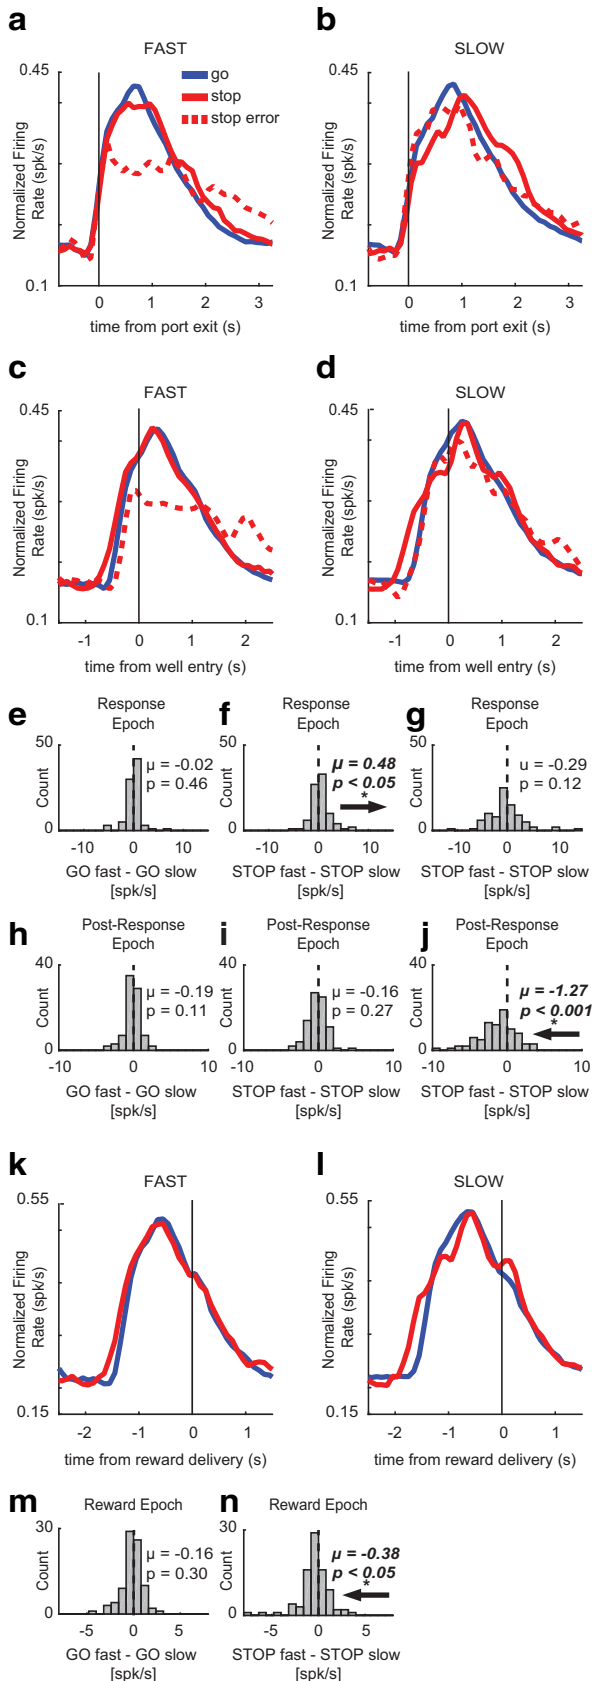
behavioral response, average population histograms were split into fast and slow trials based on movement time within each session and then averaged across sessions. ***A–D, K, L***, Average population histograms for fast and slow movement times for GO (blue), STOP (red), and STOP error (dashed) trials, aligned to port exit (***A, B***), well entry (***C, D***), and reward delivery (***K, L***). ***E–J, M, N***, Population average distributions for effects between fast and slow trials. Arrows depict direction of distribution shift for significant effects. ***A, B***, Activity was aligned to port exit for fast and slow movement times. ***E–G***, Firing rates were compared between fast and slow trial types during the response epoch by calculating trial type indices for GO trials (GO fast – GO slow), STOP trials (STOP fast – STOP slow), and STOP errors (STOP error fast – STOP error slow). During the response epoch, firing was significantly less on slow STOP trials compared to fast STOP trials (***F***), but there was no difference in firing between fast and slow GO trials (***E***) or STOP errors (***G***). ***C, D***, Activity was aligned to well entry for fast and slow movement times. ***H–J***, Firing rates were compared between fast and slow trial types during the post-response epoch by calculating trial type indices described above. During the post-response epoch, firing on fast STOP errors was significantly lower than firing on slow STOP errors (***J***). There was no significant difference in firing between fast and slow GO (***H***) or STOP trials (***I***). ***K, L***, Activity was aligned to reward delivery for fast and slow movement times. ***M, N***, Firing rates were compared between fast and slow trials during the reward epoch by calculating trial type indices for GO trials (GO fast – GO slow) and STOP trials (STOP fast – STOP slow). During the reward epoch, firing was significantly higher on slow STOP trials compared to fast STOP trials (***N***). There was no significant difference between fast and slow GO trials at the time of reward (***M***).

became worse with more preceding GO trials ($R^2 = 0.929$; $p = 0.002$; Fig. 4*F*). Parallel to this result, we found significant reductions in firing as the number of previous GO trials increased. For example, Figure 4*B,D* illustrates firing on trials in which rats performed five GO trials before a successful STOP trial (5gS). Firing was significantly reduced on 5gS trials compared to 1gS trials (Wilcoxon, $\mu = -0.12$, $p < 0.05$, df = 84).

To further quantify this effect, we computed indices to compare firing rates on sS trials and STOP trials preceded by multiple GO trials (gS – sS/gS + sS) ranging from one to six GO trials. We found a significant effect for the number of previous GO trials on the firing rate of the current STOP trial during the post-response epoch, where firing rate on STOP trials became lower as the number of preceding GO trials increased ($R^2 = 0.878$; $p = 0.006$; Fig. 4*E*). Overall, these results suggest a positive relationship between performance and DA firing such that the worse rats were on gS trials, the lower firing should be. Indeed, we found a positive correlation between the two ($R^2 = 0.69$; $p = 0.04$), demonstrating that lower probabilities of success were accompanied by reduced DA firing.

Non-DA neurons fire more on STOP trials during the response period

To determine whether firing patterns observed above were unique to putative DA neurons in VTA, we identically analyzed the 475 cells that were categorized as non-DA.
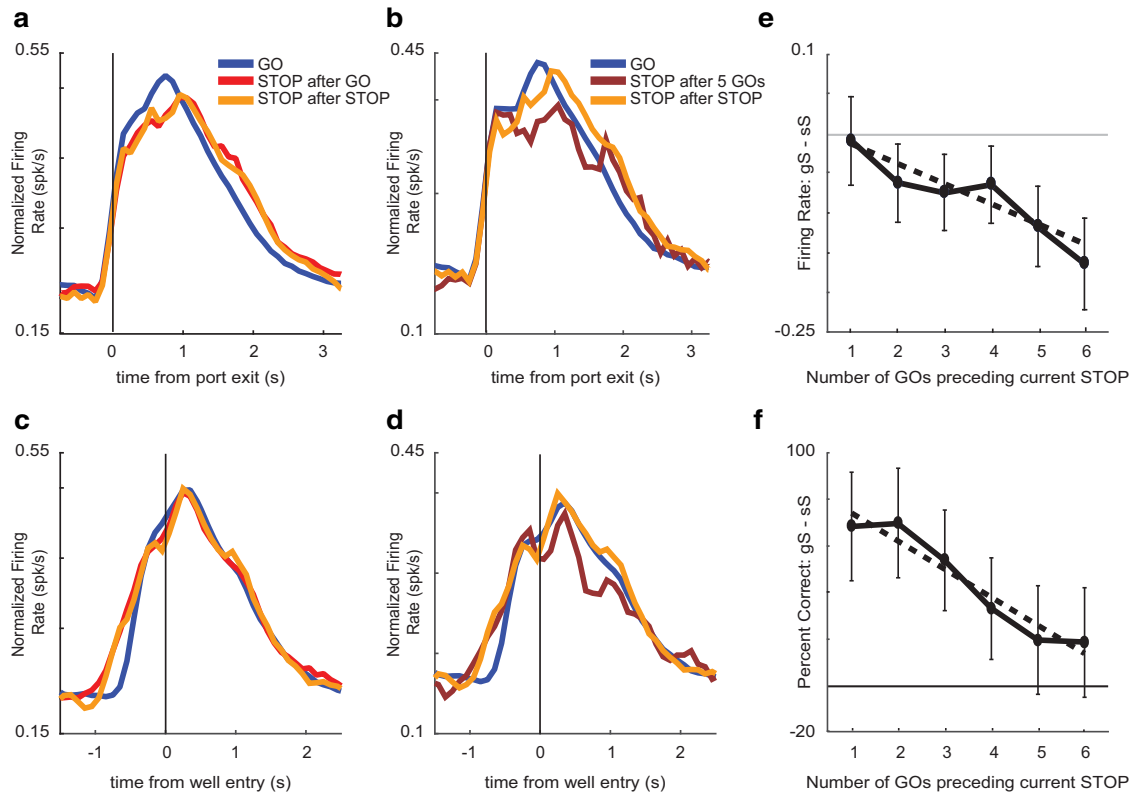
Figure 3. Movement time effects on putative DA firing. To determine whether putative DA activity was affected by speed of

Figure 4. Effects of trial order on putative DA firing. Average putative DA activity ($n = 85$) plotted for GO trials, STOP trials preceded by a GO trial (gS), and STOP trials preceded by a STOP trial (sS). For reference, GO trials have the lowest conflict, while gS trials have the highest conflict. **A, C,** Average population histogram for GO trials (blue), gS (red), and sS (orange) trial types aligned to port exit (**A**) and well entry (**C**). **B, D,** Average population histogram for GO trials (blue), STOP trials after five GO trials (maroon), and sS trials (orange) aligned to port exit (**B**) and well entry (**D**). **E,** To quantify effects, we computed indices on firing rates to compare sS to STOP trials with one to six previous GO trials (gS – sS/gS + sS) during the post-response epoch. **F,** Percentage correct behavior on STOP trials preceded by one through six GO trials.

Average firing over trial time aligned to multiple events is illustrated in Figure 5A–D. Firing of these neurons decreased on port entry and increased slightly at the time of GO cue presentation (Fig. 5A). Subsequently, after port exit, firing decreased on GO trials, but maintained a constant rate on STOP trials (Fig. 5B). Differences between trial types were not present when firing was aligned to well entry (Fig. 5C). On correct trials, firing remained low until the rat consumed the reward and exited the fluid well (data not shown). On error trials (dashed) firing decreased briefly and returned to baseline levels, again, at the time rats exited the fluid well (Fig. 5C, red dashed).

As above, we computed a trial type index (STOP – GO/STOP + GO) for firing rates taken from port exit to well entry (response-epoch) on correct trials to determine whether firing rate was significantly different between STOP and GO trials across the entire population (Wilcoxon, $p < 0.05$). The distribution was significantly shifted above zero (Wilcoxon, $\mu = 0.007$; $p < 0.005$; df = 474; Fig. 5E) and the counts of neurons that fired significantly more on STOP trials outnumbered those with the opposite effect (black bars; 45 STOP vs five GO; $\chi^2 = 31.84$; $p = 1.67^{-8}$; Fig. 5E), suggesting that firing was stronger on STOP versus GO trials during the response epoch at the population and single unit level. To quantify the dif-

ference between correct and error STOP trials, we generated the error index (STOP error – STOP correct/STOP error + STOP correct) on firing rates taken from the 800 ms period following well entry. During the post-response epoch, the distribution was significantly shifted above zero (Wilcoxon, $\mu = 0.029$; $p < 0.0001$; Fig. 5F) and the counts of neurons that fired significantly less on STOP correct trials outnumbered those that fired more on STOP error trials (63 STOP error vs 19 STOP correct; $\chi^2 = 23.50$; $p = 1.25^{-6}$; Fig. 5F, black bars). These results demonstrate that non-DA neurons fire more on STOP errors compared to correct STOP trials following well entry.

Lastly, we determined if firing of non-DA neurons would differ between GO and STOP trials at the time of reward (reward epoch; Fig. 5D) as they did for putative DA neurons. To quantify this effect, we computed the trial-type index (STOP – GO/STOP + GO) on firing rates taken during the 500 ms period following reward delivery (reward epoch). Non-DA neurons displayed no significant difference in firing between STOP and GO trials (Wilcoxon, $\mu = -0.002$; $p = 0.64$; Fig. 5G) and the counts of neurons that fired more to rewards delivered on STOP trials did not differ significantly from the counts of neurons that fired more to rewards delivered on GO trials (15 STOP vs 9 GO; $\chi^2 = 1.45$; $p = 0.23$; Fig. 5G, black bars). Overall,
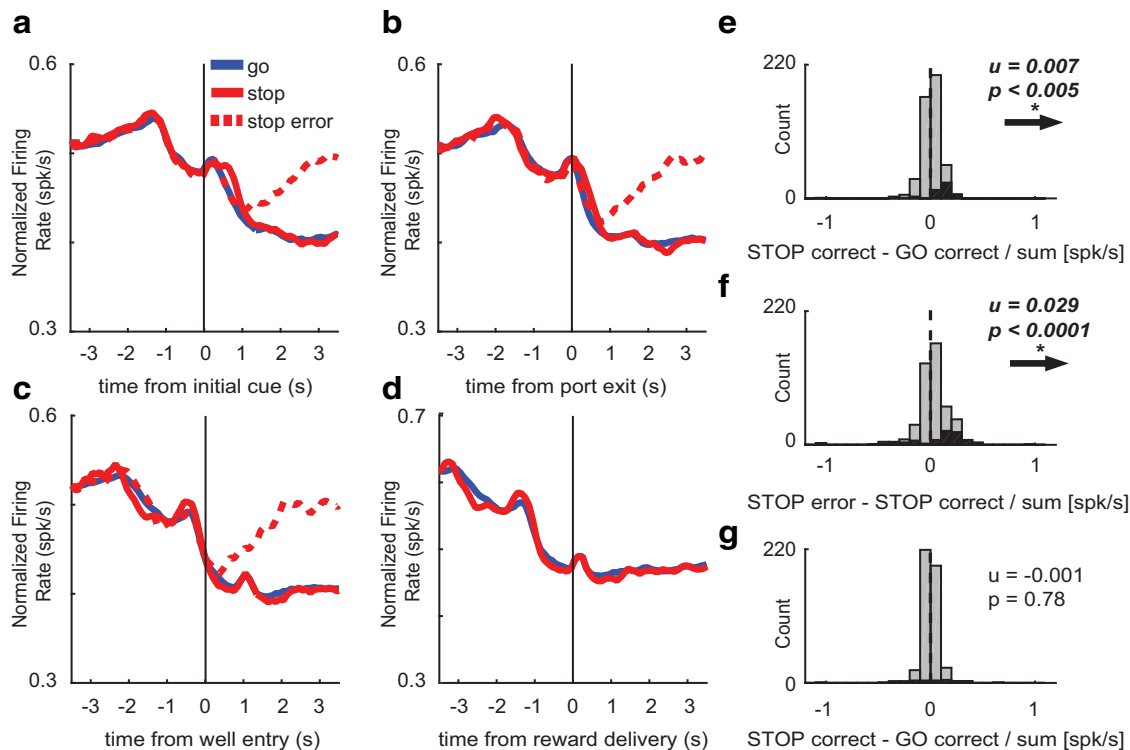
Figure 5. Characterization of non-DA firing. *A–D*, Population average histograms of non-DA neurons plotted over trial time for GO (blue), STOP (red), and STOP error (dashed) trials while aligning to multiple events (*n* = 475). Activity is aligned to initial GO cue (*A*), port exit (*B*), well entry (*C*), and reward delivery (*D*). *A*, At the time of the initial cue, firing rate increases for all trial types. *B*, During the response epoch, firing was significantly higher on STOP trials compared to GO trials. *C*, During the post-response epoch, firing was significantly higher on STOP errors compared to STOP corrects. *D*, During the reward epoch, there was no significant difference in firing between STOP and GO trials. Error trials are excluded since reward is not delivered on errors. *E–G*, Population average distributions for significant effects described above. Arrows depict direction of distribution shift for significant effects. Black bars represent the number of neurons that showed a significant difference between GO and STOP corrects (*E*, *G*) and correct and incorrect STOP trials (*F*; Wilcoxon, *p* < 0.05). Distributions are determined to be significantly different from zero via Wilcoxon. *E*, *G*, Firing rate during the response and reward epoch were compared during correct STOPs and GOs by computing a trial type index (STOP correct − GO correct/STOP correct + GO correct). *F*, Firing rate during the post-response epoch was compared in correct and incorrect STOP trials (STOP error − STOP correct/STOP error + STOP correct).

these results demonstrate that non-DA cells fire opposite to that of putative DA cells during response and post-response epochs, and do not fire more on STOP trials during reward delivery as observed for putative DA neurons.

Non-DA neuron firing is not modulated by movement speed

Next, we examined whether the speed of responding modulated the firing rate of non-DA neurons. As before, we generated average population histograms for both fast and slow trials based on movement time and examined differences in firing across three behavioral epochs (response epoch, post-response epoch, and reward epoch). To determine whether firing rate significantly differed between fast and slow trials, we calculated speed indices for firing rates on fast and slow GO trials (GO fast − GO slow), fast and slow STOP trials (STOP fast − STOP slow), and fast and slow STOP errors (fast STOP error − slow STOP error). Unlike the putative DA neurons, which fired differently on fast and slow STOP trials, non-DA neurons did

not fire differently on fast and slow GOs, STOPs, or STOP errors during any epoch.

During the response epoch, we observed no apparent differences between the firing of non-DA cells on fast and slow trial types (Fig. 6*A*,*B*). We found no significant differences in firing between fast and slow GO trials (Wilcoxon, $\mu = 0.24$; $p = 0.11$; df = 474; Fig. 6*E*), fast and slow STOP trials (Wilcoxon, $\mu = 0.00$; $p = 0.56$; df = 474; Fig. 6*F*), or fast and slow STOP error trials (Wilcoxon, $\mu = -0.08$; $p = 0.35$; df = 474; Fig. 6*G*). During the post-response epoch, firing appeared to be the same for fast and slow trial types (Fig. 6*C*,*D*). There were no significant differences between fast and slow GO trials (Wilcoxon, $\mu = -0.01$; $p = 0.99$; df = 474; Fig. 6*H*), fast and slow STOP trials (Wilcoxon, $\mu = 0.03$; $p = 0.60$; df = 474; Fig. 6*I*), or fast and slow STOP error trials (Wilcoxon, $\mu = 0.08$; $p = 0.06$; df = 474; Fig. 6*J*) during the post-response epoch. Lastly, during the reward epoch, firing appeared to be the same between fast and slow GO and STOP trials (Fig. 6*K*,*L*). There was no significant difference between fast and slow GO trials (Wilcoxon, $\mu = 0.07$; $p = 0.46$; Fig. 6*M*), or fast and
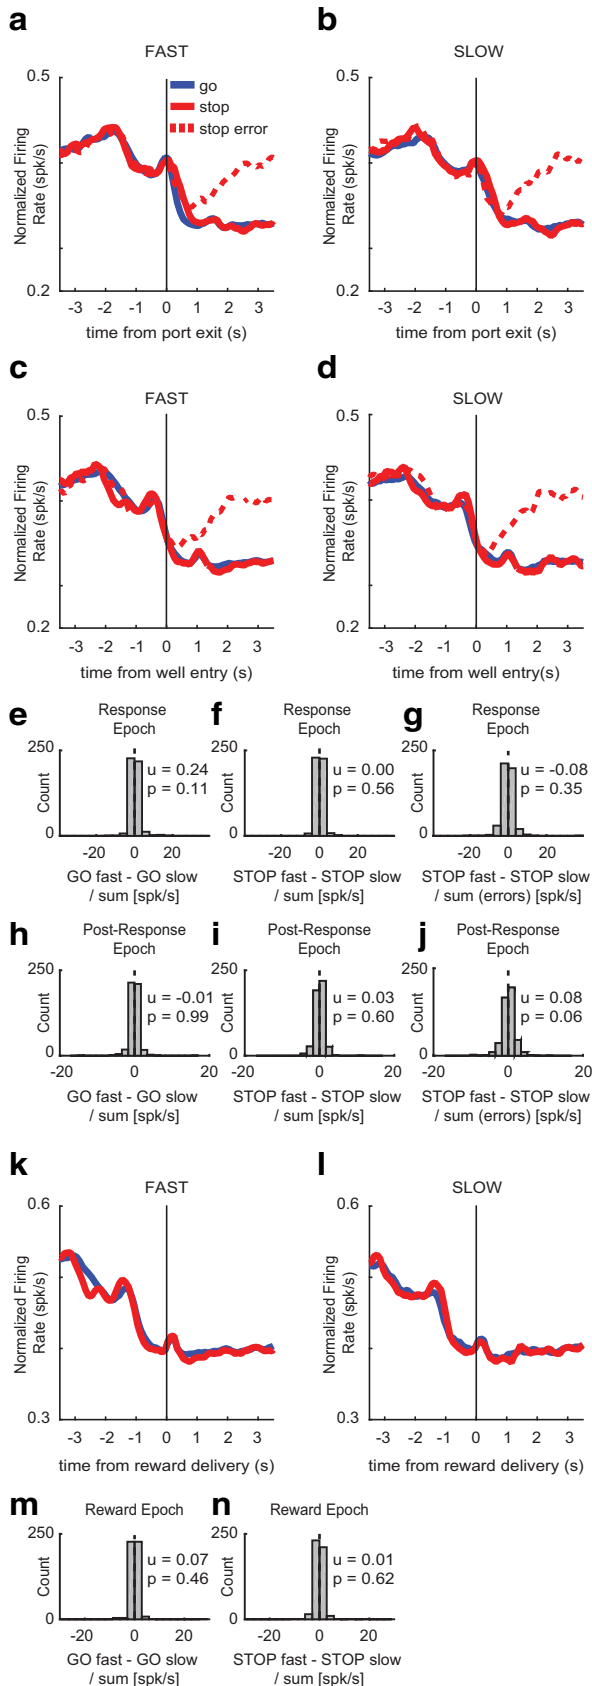
response, average population histograms were split into fast and slow trials based on movement time. **A–D, K, L**, Average population histograms for fast and slow movement times for GO (blue), STOP (red), and STOP error (dashed) trials, aligned to port exit (**A, B**), well entry (**C, D**), and reward delivery (**K, L**). **A, B**, Activity was aligned to port exit for fast and slow movement times. **E–G**, Firing rates were compared between fast and slow trial types during the response epoch by calculating trial type indices for GO trials (GO fast – GO slow/GO fast + GO slow), STOP trials (STOP fast – STOP slow/STOP fast + STOP slow), and STOP errors [STOP fast – STOP slow/STOP fast + STOP slow (errors)]. **C, D**, Activity aligned to well entry for fast and slow movement times. **H–J**, Firing rates were compared between fast and slow trial types during the post-response epoch using indices described above. **K, L**, Activity was aligned to reward delivery for fast and slow movement times. **M, N**, Firing rates were compared between fast and slow trials during the reward epoch by calculating the same trial-type indices described above.

slow STOP trials (Wilcoxon, $\mu = 0.01$; $p = 0.62$; df = 474; Fig. 6*N*) at the time of reward.

Non-DA neuron firing was modulated by heightened response conflict induced by previous trial type

Lastly, we investigated whether the previous trial type would affect non-DA firing on the current trial. To determine whether non-DA cells might contribute to behavior modulation based on previous trial type, we compared average firing activity on correct GO trials, STOP trials preceded by up to six preceding GO trials (gS), and STOP trials following a STOP trial (sS). A regression analysis comparing percentage correct between the six trial types revealed a significant effect between the number of previous GO trials and accuracy on the current STOP trial during sessions during which non-DA neurons were recorded ($R^2 = 0.898$; $p = 0.004$; Fig. 7*F*). We also found that firing on current STOP trials was modulated by the number of preceding GO trials. As an example, we plotted firing on STOP trials preceded by five GO trials (Fig. 7*B,D*). As before, this effect was quantified by computing indices to compare firing rates between sS trials and STOP trials preceded by multiple GO trials (gS – sS/gS + sS) during the post-response epoch (Fig. 7*E*). During the post-response epoch, we found a significant effect for the number of previous GO trials on firing rate of the current STOP trial, such that firing rate on STOP trials became lower as the number of preceding GO trials increased ($R^2 = 0.887$; $p = 0.005$).

Discussion

In this study, we recorded activity of putative DA neurons in the VTA from rats performing a stop-change task. Phasic bursts in DA have been shown to reflect both reward prediction errors and saliency in a variety of tasks, yet it is unclear how DA firing is modulated as a function of performance on tasks that require response inhibition and cognitive control. On the one hand, if DA activity reflects changes in value associated with a particular cue, we would expect to see decreased DA firing at the time of cues that predict a lower probability of receiving reward
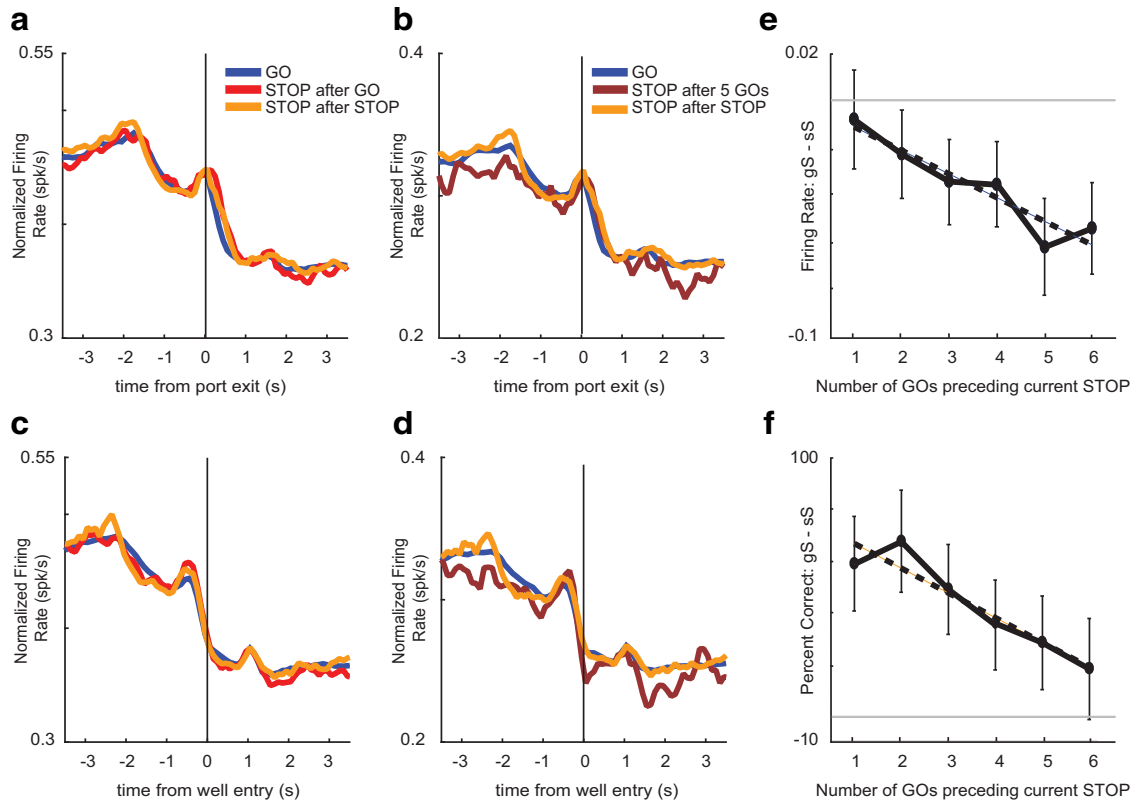
Figure 6. Movement time effects on non-DA firing. To determine whether non-DA activity was affected by speed of behavioral

Figure 7. Effects of trial order on non-DA firing. Average non-DA activity (n = 475) plotted for GO trials, STOP trials preceded by a GO trial (gS), and STOP trials preceded by a STOP trial (sS). **A**, **C**, Average population histogram for GO trials (blue), gS (red), and sS (orange) trial types aligned to port exit (**A**) and well entry (**C**). **B**, **D**, Average population histogram for GO trials (blue), STOP trials after five GO trials (maroon), and sS trials (orange) aligned to port exit (**B**) and well entry (**D**). **E**, To quantify effects, we computed indices on firing rates to compare sS to STOP trials with one to six previous GO trials gS to GO trials (gS – GO/gS + GO) during the post-response epoch. **F**, Percentage correct behavior on STOP trials preceded by one through six GO trials.

(i.e., STOP cues), followed by elevated DA firing at the time of unexpected reward. In this light, DA activity can be thought of as an indirect indicator of an animal's intuition about its probability of success on a given trial. In support of this hypothesis, previous literature has shown that DA neurons encode reward probabilities and outcomes, such that cues predicting a lower probability of reward yield weaker phasic DA responses, while unexpected and low probability rewards yield phasic increases in DA activity (Fiorillo et al., 2003). On the other hand, if DA activity reflects saliency or the need to inhibit behavior, we would expect to see increased firing of DA neurons to STOP cues.

We found very few DA neurons that fired significantly stronger on STOP trials during the response epoch, and that firing to STOP cues was dependent on the identity of the previous trial type, in agreement with the behavioral evidence for heightened response conflict. We found that across the population and at the single neuron level, putative DA neurons exhibited lower and higher firing to STOP cues and rewards, respectively. This prediction error effect was enhanced as the number of preceding GO trials increased, such that DA's activity was modulated by the conflict associated with an unexpected STOP trial following a train of GO trials. These findings support our idea that DA signaling on the stop-change task is indica-

tive of an animal's sense about its future probability of success, and are also in line with previously published work linking the activity of midbrain DA neurons to an animal's belief in choice accuracy during a perceptual decision-making task (Lak et al., 2017). In that study, using a computational modeling approach, DA neurons were shown to be sensitive to reward prediction error and the same signal also represented statistical certainty in reward (Lak et al., 2017).

In many ways, these findings are supported by other studies that suggest two distinct populations of DA neurons exist (Matsumoto and Hikosaka, 2009). Recordings from substantia nigra pars compacta (SNc) reveal a gradient in the density of neurons encoding motivational value versus saliency with greater numbers of value-encoding neurons found along the ventromedial extent (Matsumoto and Hikosaka, 2009). This is somewhat in contrast with the VTA, where high numbers of value-encoding neurons have been described, along with sparse numbers of saliency encoding neurons (Matsumoto and Hikosaka 2009, Bromberg-Martin et al., 2010). These findings, originally described in non-human primates, largely fit with the results we present here, and suggests that DA neurons in SNc might report the salience of stop cue during performance of our stop-change task.

Previously, pharmacologic control of response inhibition via the DA system has yielded conflicting results on DA's role in GO and STOP trial performance (Eagle and Baunez, 2010). In children and adults with ADHD, administration of methylphenidate and D-amphetamine, psychostimulants that target the DA system, have reported improvement in the stop signal reaction time (SSRT), a common measure of response inhibition in which the longer the SSRT the more time the animal needs to inhibit the response (Tannock et al., 1989; Aron et al., 2003; Boonstra et al., 2005). However, other studies using the same psychostimulants in hyperactive children report improvements on GO reaction time but not SSRT, as well as other forms of unwanted impulsivity (Bedard et al., 2003; Lijffijt et al., 2006). Still, there is evidence that DA's effects may be baseline dependent, decreasing SSRT in slow responders, and increasing or worsening SSRT, in subjects with fast SSRTs (Boonstra et al., 2005; Eagle et al., 2007). Collectively, these results suggest that systemic increases in DA may non-selectively enhance either GO or STOP performance, which may be dependent on the subject's baseline performance at the start of the experiment. DA's role in action control may be better explained at the receptor level. A study that administered GBR-12909, a DA reuptake inhibitor, reported no effect on SSRT in rats (Bari et al., 2009). Administration of cis-flupenthixol, a mixed D1/D2 receptor antagonist, also had no impact on SSRT in rats, and failed to block the SSRT-decreasing effect of methylphenidate (Eagle et al., 2007). However, administration of either a D1- or D2-receptor antagonist (SCH23390 or sulpiride) directly into the dorsal medial striatum (DMS) led to opposing effects on SSRT, with SSRT decreasing after D1 antagonism, and increasing after D2 antagonism (Eagle and Baunez, 2010). While the DMS is well-characterized regarding its role in decision-making and action selection and initiation, further research is needed to assess how the DMS decodes midbrain DA input in the context of the stop-signal task.

The nucleus accumbens (NAc) also receives DA projections and plays a key role in reward-seeking and impulsivity. In one study, the depletion of DA in the NAc attenuated amphetamine-induced increases in premature responding during a 5-choice serial reaction time task (5-CSRTT; Cole and Robbins, 1989). Remarkably, amphetamine-induced increases in premature responding could be blocked by systemic administration of D1/D2 mixed receptor antagonist, cis-flupenthixol, and D1 receptor antagonist, SCH23390, although these antagonists had no impact on SSRT (Eagle and Baunez, 2010). Increased impulsivity is also associated with reduced D2/D3 receptor activity in the NAc (Eagle and Baunez, 2010; Dalley and Robbins, 2017). These findings suggest that NAc DA transmission may modulate impulsivity via receptor-mediated processes. There is also evidence suggestingt that NAc DA release is modulated by action initiation. In a study that measured DA concentration in freely-behaving rats during a go/no-go task, NAc DA increased on no-go trials only after correct movement was initiated for this trial type (Syed et al., 2016). These findings are suggestivet that downstream consequences of DA in higher level processing areas are critical for task performance.

The role of DA in the regulation of impulsivity is largely dependent on the type of impulsivity being tested (Dalley and Robbins, 2017). Waiting impulsivity, an animal's ability to refrain from responding until receiving a specific cue or amount of time has elapsed, is thought to be influenced by the activity of DA cells arising from the VTA that project to the ventral striatum. However, response inhibition or "stopping impulsivity," an animal's ability to stop and redirect a prepotent action, is largely dependent on the action of DA in the dorsal striatum (Dalley and Robbins, 2017). It is unclear to what degree DA from VTA neurons influence or support response inhibition. Since DA neurons in VTA strongly project to NAc, we suspect the role that VTA DA neurons have on stop-signal performance is to track the probability of reward, as opposed to signaling saliency or the need to inhibit behavior. As suggested above, if we were to record from SNc that contains DA neurons projecting strongly to dorsal striatum, we might observe a higher percentage of neurons that fire more strongly to STOP cues, thus playing a more direct role in response inhibition.

Critically, we do not see similar changes in non-DA neurons. In the main analyses we found that non-DA cells in the VTA actually showed the opposite effect; higher firing on STOP trials compared to GO trials. However, this effect might simply reflect differences in when the movement was terminated, because firing did not increase above GO-induced firing but simply extended until the end of the movement. Non-DA cells did, however, exhibit sequence effects that mirrored those seen in DA cells such that increased numbers of GO trials before a STOP trial resulted in reduced activity and a decreased likelihood of responding correctly. These signals may reflect reward expectancy signals either in the form of future motor planning events or being reflective of the overall integration events occurring in VTA (Watabe-Uchida et al., 2017).

The stop-change task requires animals to inhibit a GO response in the presence of a STOP cue and the realization and utilization of this strategy is integral to successful performance. In the goal-directed behavior literature, there is new debate surrounding the involvement of DA in the interplay between model-based and model-free behaviors (Langdon et al., 2018). Phasic bursts in DA have traditionally been interpreted as model-free generated prediction errors as an animal encounters valuable information or reward (Langdon et al., 2018). In our task, increases in DA to reward on STOP trials could reflect general learning or "surprise" associated with receiving reward. However, given that these increases in DA firing to reward also occurred on trials in which rats adaptively slowed their behavior, these changes in DA firing may be reflective of the rats adopting a model-based approach. There is reason to think that DA signals are heterogeneous, and not simply scalar representations of value independent of the form of the expected reward (Sadacca

et al., 2016; Soares et al., 2016; Starkweather et al., 2017; Langdon et al., 2018). If true, components of the overall DA response may reflect model-free and model-based predictions. Future research should explore the specific temporal components of this response to further elucidate the exact computational support VTA DA cells are offering neural networks supporting response inhibition.

Response inhibition is a complex and dynamic behavior that is reliant on several brain regions. In our study, we show that DA signaling in VTA neurons appears to reflect the uncertainty associated with a low probability of reward on STOP trials. This is distinct from commonly posited beliefs that VTA DA signaling is necessary to engage in response inhibition or provides a neural correlate of saliency associated with the low occurrence of STOP trials. These data are some of the first to characterize a function for VTA DA neurons during a task that requires response inhibition.

References

Aron AR, Dowson JH, Sahakian BJ, Robbins TW (2003) Methylphenidate improves response inhibition in adults with attention-deficit/hyperactivity disorder. Biol Psychiatry 54:1465–1468. Medline

Bari A, Eagle DM, Mar AC, Robinson ES, Robbins TW (2009) Dissociable effects of noradrenaline, dopamine, and serotonin uptake blockade on stop task performance in rats. Psychopharmacology 205:273–283. CrossRef Medline

Bedard AC, Ickowicz A, Logan GD, Hogg-Johnson S, Schachar R, Tannock R (2003) Selective inhibition in children with attention-deficit hyperactivity disorder off and on stimulant medication. J Abnorm Child Psychol 31:315–327. Medline

Bellgrove MA, Chambers CD, Vance A, Hall N, Karamitsios M, Bradshaw JL (2006) Lateralized deficit of response inhibition in early-onset schizophrenia. Psychol Med 36:495–505. CrossRef Medline

Boecker M, Gauggel S, Drueke B (2013) Stop or stop-change - Does it make any difference for the inhibition process? Int. J. Psychophysiol 87:234–243. CrossRef

Boonstra AM, Kooij JJ, Oosterlaan J, Sergeant JA, Buitelaar JK (2005) Does methylphenidate improve inhibition and other cognitive abilities in adults with childhood-onset ADHD? J Clin Exp Neuropsychol 27:278–298. CrossRef Medline

Bromberg-Martin ES, Matsumoto M, Hikosaka O (2010) Dopamine in motivational control: rewarding, aversive, and alerting. Neuron 68:815–834. CrossRef

Bryden DW, Burton, Kashtelyan V, Barnett BR, Roesch MR (2012) Response inhibition signals and miscoding of direction in dorsomedial striatum. Front Integr Neurosci 6:69.

Bryden DW, Roesch MR (2015) Executive control signals in orbitofrontal cortex during response inhibition. J Neurosci 35:3903–3914. CrossRef Medline

Bryden DW, Johnson EE, Tobia SC, Kashtelyan V, Roesch MR (2011) Attention for learning signals in anterior cingulate cortex. J Neurosci 31:18266–18274. CrossRef Medline

Bryden DW, Brockett AT, Blume E, Heatley K, Zhao A, Roesch MR (2018) Single neurons in anterior cingulate cortex signal the need to change action during performance of a stop-change task that induces response competition. Cereb Cortex Advance online publication, February 3, 2018. CrossRef

Cole BJ, Robbins TW (1989) Effects of 6-hydroxydopamine lesions of the nucleus accumbens septi on performance of a 5-choice serial reaction time task in rats: implications for theories of selective attention and arousal. Behav Brain Res 33:165–179. CrossRef

Dajani DR, Uddin LQ (2015) Demystifying cognitive flexibility: implications for clinical and developmental neuroscience. Trends Neurosci 38:571–578. CrossRef Medline

Dalley JW, Robbins TW (2017) Fractionating impulsivity: neuropsychiatric implications. Nat Rev Neurosci 18:158–171. CrossRef Medline

Durston S, de Zeeuw P, Staal WG (2009) Imaging genetics in ADHD: a focus on cognitive control. Neurosci Biobehav Rev 33:674–689. CrossRef Medline

Eagle DM, Baunez C (2010) Is there an inhibitory-response-control system in the rat? Evidence from anatomical and pharmacological studies of behavioral inhibition. Neurosci Biobehav Rev 34:50–72. CrossRef Medline

Eagle DM, Tufft MR, Goodchild HL, Robbins TW (2007) Differential effects of modafinil and methylphenidate on stop-signal reaction time task performance in the rat, and interactions with the dopamine antagonist cis-flupenthixol. Psychopharmacology 192:193–206. CrossRef Medline

Eagle DM, Bari A, Robbins TW (2008) The neuropsychopharmacology of action inhibition: cross species translation of the stop-signal and go/no-go tasks. Psychopharmacology 199:439–456. CrossRef Medline

Fillmore MT, Rush CR (2002) Impaired inhibitory control of behavior in chronic cocaine users. Drug Alcohol Depend 66:265–273. Medline

Fiorillo CD, Tobler PN, Schultz W (2003) Discrete coding of reward probability and uncertainty by dopamine neurons. Science 299:1898–1902. CrossRef Medline

Lak A, Nomoto K, Keramati M, Sakagami M, Kepecs A (2017) Midbrain dopamine neurons signal belief in choice accuracy during a perceptual decision. Curr Biol 27:821–832. CrossRef

Langdon AJ, Sharpe MJ, Schoenbaum G, Niv Y (2018) Model-based predictions for dopamine. Curr Opin Neurobiol 49:1–7. CrossRef Medline

Lijffijt M, Kenemans JL, ter Wal A, Kemner C, Westenberg H, Verbaten MN, van Engeland H (2006) Dose-related effect of methylphenidate on stopping and changing in children with attention-deficit/hyperactivity disorder. Eur Psychiatry 21:544–547. CrossRef

Jo YS, Lee J, Mizumori SJY (2013) Effects of prefrontal cortical inactivation on neural activity in the ventral tegmental area. J Neurosci 33:8159–8171. CrossRef

Matsumoto M, Hikosaka O (2009) Two types of dopamine neuron distinctly convey positive and negative motivational signals. Nature 459:837–841. CrossRef Medline

Monterosso JR, Aron AR, Cordova X, Xu J, London ED (2005) Deficits in response inhibition associated with chronic methamphetamine abuse. Drug Alcohol Depend 79:273–277. CrossRef Medline

Oosterlaan J, Sergeant JA (1998) Response inhibition and response re-engagement in attention-deficit/hyperactivity disorder, disruptive, anxious and normal children. Behav Brain Res 94:33–43. CrossRef

Oosterlaan J, Logan GD, Sergeant JA (1998) Response inhibition in AD/HD, CD, comorbid AD/HD + CD, anxious, and control children: a meta-analysis of studies with the stop task. J Child Psychol Psychiatry 39:411–425. CrossRef

Park J, Moghaddam B (2017) Risk of punishment influences discrete and coordinated encoding of reward- guided actions by prefrontal cortex and VTA neurons. Elife 6:e30056. CrossRef

Paxinos G, Watson C (2007) The rat brain in stereotaxic coordinates. London: Academic Press.

Roesch MR, Taylor AR, Schoenbaum G (2006) Encoding of time-discounted rewards in orbitofrontal cortex is independent of value representation. Neuron 51:509–520. CrossRef

Roesch MR, Calu DJ, Schoenbaum G (2007) Dopamine neurons encode the better option in rats deciding between differently delayed or sized rewards. Nat Neurosci 10:1615–1624. CrossRef

Sadacca B, Jones JL, Schoenbaum G (2016) Midbrain dopamine neurons compute inferred and cached value predictions errors in a common framework. Elife 5:e13665. CrossRef

Schultz W (2013) Updating dopamine reward signals. Curr Opin Neurobiol 23:229–238. CrossRef

Soares S, Atallah BV, Paton JJ (2016) Midbrain dopamine neurons control judgment of time. Science 354:1273–1277. CrossRef Medline

Starkweather CK, Babayan BM, Uchida N, Gershman SJ (2017) Dopamine reward prediction errors reflect hidden-state inference across time. Nat Neurosci 20:581–589. CrossRef Medline

Syed EC, Grima LL, Magill PJ, Bogacz R, Brown P, Walton ME (2016) Action initiation shapes mesolimbic dopamine encoding of future rewards. Nat Neurosci 19:34–36. CrossRef

Takahashi YK, Schoenbaum G (2016) Ventral striatal lesions disrupt dopamine neuron signaling of differences in cue value caused by changes in reward timing but not number. Behav Neurosci 130: 593–599. CrossRef

Takahashi YK, Roesch M, Stalnaker TA, Haney RZ, Calu DJ, Taylor AR, Burke KA, Schoenbaum G (2009) The orbitofrontal cortex and ventral tegmental area are necessary for learning from unexpected outcomes. Neuron 62:269–280. CrossRef Medline

Takahashi YK, Roesch M, Wilson RC, Toreson K, O'Donnell P, Niv Y, Schoenbaum G (2011) Expectancy-related changes in firing of dopamine neurons depend on orbitofrontal cortex. Nat Neurosci 14:1590–1597. CrossRef

Takahashi YK, Langdon AJ, Niv Y, Schoenbaum G (2016) Temporal specificity of reward prediction errors signaled by putative dopamine neurons in rat VTA depends on ventral striatum. Neuron 91:182–193. CrossRef Medline

Tannock R, Schachar RJ, Carr RP, Chajczyk D, Logan GD (1989) Effects of methylphenidate on inhibitory control in hyperactive children. J Abnorm Child Psychol 17:473–491. Medline

Verbruggen F, Logan GD (2008) Response inhibition in the stop-signal paradigm. Trends Cogn. Sci 12:418–424. CrossRef Medline

Watabe-Uchida M, Eshel N, Uchida N (2017) Neural circuitry of reward prediction error. Annu Rev Neurosci 40:373–394. CrossRef

Wood J, Simon NW, Koerner FS, Kass RE, Moghaddam B (2017) Networks of VTA neurons encode real-time information about uncertain numbers of actions executed to earn a reward. Front Behav Neurosci 11:140.